

Pricing Algorithms and Tacit Collusion

Bruno Salcedo*

Pennsylvania State University

This version — 11 · 1 · 2015

[Click here for current version](#)

Abstract There is an increasing tendency for firms to use pricing algorithms that speedily react to market conditions, such as the ones used by major airlines and online retailers like Amazon. I consider a dynamic model in which firms commit to pricing algorithms in the short run. Over time, their algorithms can be revealed to their competitors and firms can revise them, if they so wish. I show how pricing algorithms not only facilitate collusion but inevitably lead to it. To be precise, within certain parameter ranges, in *any* equilibrium of the dynamic game with algorithmic pricing, the joint profits of the firms are close to those of a monopolist.

Keywords Algorithmic pricing · Tacit collusion · Repeated games

JEL classification L13 · D43 · C73 · K21

*brunosalcedo.com · bruno@psu.edu · This paper was written as part of my dissertation under the invaluable guidance and supervision of Edward Green, Vijay Krishna and Nageeb Ali. For helpful comments and discussions, I am grateful to Ruilin Zhao, Kalyan Chatterjee, Rohit Lamba, Neil Wallace, Henrique De Oliveira, Ron Siegel, Bruno Sultanum, Evelyn Nunes, Nail Kashaev, Gustavo Gudiño, and David Rahman, as well as seminar attendants at The Pennsylvania State University, and the PSU-Cornell 2015 Fall Macroeconomics Conference. I gratefully acknowledge the [Human Capital Foundation](#) and particularly Andrey P. Vavilov, for research support through the [Center for the Study of Auctions, Procurements, and Competition Policy](#) at the Pennsylvania State University. Any remaining errors are my own.

1. Introduction

In the Spring of 2011, two online retailers offered copies of Peter Lawrence’s textbook *The Making of a Fly* on Amazon for \$18,651,718.08 and \$23,698,655.93 respectively. This was the result of both sellers using automated *pricing algorithms*. Everyday, the algorithm used by seller 1 set the price of the book to be 0.9983 times the price charged by seller 2. Later in the day, seller 2’s algorithm would adjust its price to be 1.27059 times that of seller 1. Prices increased exponentially and remained over one million dollars for at least ten days (!), until one of the sellers took notice and adjusted its price to \$106.23.¹

Automated pricing algorithms—presumably better than the ones outlined in the previous story—are now ubiquitous in many different industries including airlines (Borenstein, 2004), online retail (Ezrachi and Stucke, 2015) and high-frequency trading (Boehmer et al., 2015). Also, while not necessarily automated, algorithms also feature in hierarchical firms in which top managers design protocols for lower level employees to implement. Optimal pricing algorithms can be highly profitable, as they would be sophisticated enough to recognize and take advantage of profitable collusion opportunities.

In this paper, I formalize these ideas via a model of dynamic competition in continuous time, in which two firms use algorithms to set prices. These algorithms react both to demand conditions and to rivals’ prices. At exogenous stochastic times, the current algorithm of a firm becomes apparent to the other firm—because it has either been inferred or “decoded”. When a firm has decoded its rival’s algorithm, it can revise its own algorithm in response.² It is important in my model that this decoding takes (stochastically) longer than the arrival of demand shocks. For instance, demand shocks could arrive weekly on average, while it may take a firm, again on average, up to six months to decode its rival’s algorithm and to implement a new algorithm itself. My main result is the following³

Theorem: When demand shocks occur much more frequently than algorithm revisions, the long-run profits from *any* equilibrium are close to those of a monopolist.

My main result differs sharply from our traditional understanding of dynamic competition. In most models of dynamic competition, there is a plethora of

¹See *Amazon’s \$23,698,655.93 book about flies*, <http://www.michaeleisen.org/blog/?p=358>.

²This revision structure is similar to that of revision games (Kamada and Kandori, 2009) or Calvo’s pricing model (Calvo, 1983). An important difference is that, in these papers, agents’ choices consists of a single act. In contrast, agents in my model choose algorithms that can react to market conditions in the short run even before new revision opportunities arise.

³A formal statement is contained in Section 4 as Theorem 4.1.

equilibria—some collusive and some not—specially when firms are patient. In the context of repeated games, such results are dubbed “folk theorems” (see [Mailath and Samuelson \(2006\)](#)). This lack of predictability is often viewed as a weakness of the theory.⁴ An early model in which cooperation the unique prediction is due to [Aumann and Sorin \(1989\)](#). Players in their model have limited memory—this is also a feature of all algorithms—and a small amount of incomplete information. Their uniqueness result, however, applies only to games of common interest. My model predicts cooperation in a more general class of games. The main message of this paper is that, when firms compete via algorithms that are fixed in the short run but can revised over time, collusion is not only possible but rather, it is *inevitable*.

1.1. Key strategic intuitions

The key mechanism underlying my main result can be understood by means of a simple example. Consider a situation in which both firms have adopted the rather simple algorithm which always prices competitively, say, at the level dictated by the one-shot Bertrand equilibrium. Let us see why this cannot be an equilibrium of the dynamic game in which algorithms can be revised. At the first opportunity, firm 1 could deviate to an algorithm that also prices competitively—thus best responding in the short run—but is programmed to match any price increases by its rival. Such a deviation can be thought of as a “proposal” to collude, which will be understood by firm 2 once it decodes firm 1’s algorithm. When this happens, firm 2 will understand that, by raising its own price, it can transition to a more profitable high-price regime. Moreover, suppose that it will then take some time for firm 1 to decode firm 2’s algorithm and revise its own in turn. Firm 2 would then expect the high-price regime to last long enough for the transition to be profitable. Hence, there cannot be equilibria in which firms set low prices forever. While this argument shows that competitive pricing is not an equilibrium, it does not say that equilibrium prices must be close to monopoly prices. This requires a more detailed analysis which is carried out in [Section 4](#).

The model has four key features leading to the result that collusion is inevitable. First, and foremost, is that *commitment* is feasible. Because it takes time revise an algorithm, firms are committed in the short run to a pricing pol-

⁴For example, [Green et al. \(2014\)](#) state that “folk theorems deal with the implementation of collusion, and have nothing to say about its initiation. The folk theorem itself does not address whether firms would choose to play the strategies that generate the monopoly outcome nor how firms might coordinate on those strategies.” Then, they proceed with the following quote from [Ivaldi et al. \(2003\)](#) “While economic theory provides many insights on the nature of tacitly collusive conducts, it says little on how a particular industry will or will not coordinate on a collusive equilibrium, and on which one.”

icy.⁵ If algorithms could be revised very frequently, firm 2 might not accept firm 1’s “proposal”—there is the danger that firm 1 would change its algorithm right away. Interestingly, the second key feature is that commitment, while feasible, is *imperfect*. The result would not hold if firms could never revise algorithms. With full commitment, the standard folk theorem arguments would apply. In particular, it would be an equilibrium for both firms to choose a simple algorithm that always implemented the Bertrand price. Hence, the result relies crucially on the fact that firms can revise their algorithms and hence commitment is not perfect.

Third, it is important that pricing algorithms are *responsive* to market outcomes. Without this feature, firm 1 would not be able to offer a price increase while at the same time best-responding in the short run. Finally, it is also important that pricing algorithms are directly *observable* or decodable by rivals. Otherwise, firm 2 would not understand firm 1’s “proposal” to switch regimes.

The proof of the main result mimics the argument of the previous example. First, I show that collusive “proposals” must be accepted when revision opportunities are infrequent. More precisely, after observing that its rival is using a specific kind of algorithm, a firm must play a dynamic best response to this algorithm for a long time (Lemma 4.2), which results in continuation values that are close to the Pareto frontier (Lemma 4.3). Then, I show that firms can make such “proposals” without affecting the path of play before their rival’s next revision. Consequently, after both firms have had a chance to revise their algorithms at least once, continuation values after each subsequent revision must be close to the Pareto frontier with high probability (Lemma 4.4). An additional step is needed. Since revisions are infrequent, it is important that continuation values are *sufficiently* close to the Pareto frontier to guarantee that the firm’s joint profits will remain high with high probability until a new revision arises.

1.2. Additional results

One should expect that it is in each firm’s best interest to have a transparent algorithm that can be decoded by its rival as quickly as possible. On one hand, this helps firms to coordinate on collusive outcomes. On the other, being relatively slower in decoding would mean that the firm is relatively more committed, and would enjoy greater bargaining power. I formalize these ideas in section 5. First, I consider an alternative specification in which firms have the option to obfuscate their algorithms so they can never be decoded. I show that, in equilibrium, they will never choose to obfuscate their algorithms (Proposition 5.1). Then, I consider the asymmetric case in which one firm revises its algorithm more frequently than

⁵Other than the time it takes a firm to decode its rival’s algorithm, similar forms of commitment could arise because of the cost or time it takes to design and implement new algorithms, or the limited attention of the people in charge of doing so.

the other. In the limit, when one firm is extremely committed and the other one can revise arbitrarily often, there is a unique equilibrium outcome in which the committed firm acts as an Stackelberg leader that extracts all the market profits (Proposition 5.2).

In section 5.3, I consider yet another means by which pricing algorithms can foster collusion, although its nature is completely different from that of the main result. Proposition 5.3, states that, even if firms are very *impatient*—and thus collusion would not be possible in a repeated game without pricing algorithms—they can sustain monopolistic profits provided that revision opportunities are frequent enough. This result arises because the ability to observe algorithms improves the firm’s monitoring technology. If algorithms can be observed and revised frequently enough, any potential deviation can be detected before any customer arrives to the market, and the revising firm can thus react to it before it becomes profitable.⁶

While the model is intended to study a duopoly, the results have general implications for the theory of repeated games. This is because none of the results depend crucially on the specific details of the stage game. The only properties of the stage game that play a crucial role in the derivation of the results are the compactness and convexity of the set of feasible profits. These properties are satisfied, for instance, in any repeated game with finite action spaces and public randomization. Similar results should obtain in different settings, as long choices are made via automated algorithms that can be decoded by rivals and cannot be revised too frequently.

1.3. Antitrust implications

My findings suggest that pricing algorithms are an effective tool for *tacit* collusion. Hence, as more firms employ pricing algorithms, there may be a need to adjust how anti-trust regulations are enforced. According to Harrington (2015),

“In the U.S., unlawful collusion has come to mean that firms have an *agreement* to coordinate their behavior... [E]videntiary standards for determining liability are based on *communications* that could produce mutual understanding and *market behavior* that is the possible consequence of mutual understanding.”

If collusion is reached through the use of pricing algorithms without explicit agreement or direct communication, it might fall outside the scope of the current regulatory framework. Indeed, according to Ezechia and Stucke (2015),

⁶Similar results arise in different models in which agents can observe information about the future plans of others, including quick-response equilibrium Anderson (1984), program equilibrium (Tennenholtz, 2004), self-referential games (Levine and Pesendorfer, 2007, Block and Levine, 2015) and revision games (Kamada and Kandori, 2009, 2011).

“Absent the presence of an agreement to change market dynamics, most competition agencies may lack enforcement tools, outside merger control, that could effectively deal with the change of market dynamics to facilitate tacit collusion through algorithms.”

See [Ezrachi and Stucke \(2015\)](#) and [Mehra \(2015\)](#) for a detailed discussion of the legal aspects and challenges of antitrust enforcement regarding industries that use pricing algorithms.

1.4. Related literature

Communication and collusion.— The role of dynamic incentives to facilitate collusion has been thoroughly studied since the seminal work of [Friedman \(1971\)](#). However, most of the literature has focused on the possibility, rather than the inevitability of collusion. [Green et al. \(2014\)](#) argue that, in many cases, there is reason to believe that communication might be necessary to *initiate* collusion.⁷ [Harrington \(2015\)](#) finds conditions on firm’s beliefs that are *sufficient* to guarantee collusive outcomes without further communication. Specifically, he assumes that it is common knowledge among firms that “any price increase will be at least matched by the other firms, and failure to do so results in the competitive outcome”. My model suggests that such common understanding could arise naturally when firms set prices through algorithms that can be decoded over time and cannot be continuously revised.

Renegotiation.— The mechanics of my model has some resemblance with the literature on renegotiation. Early work on *implicit* renegotiation imposed axiomatic restrictions on the set of equilibria based on the idea that different equilibria cannot be Pareto ranked, because players would always renegotiate and opt for the Pareto dominant alternative ([Pearce, 1987](#), [Bernheim and Ray, 1989](#), [Farrell and Maskin, 1989](#)). More recent work has focused on explicit renegotiation protocols.

[Miller and Watson \(2013\)](#) consider a model of explicit bargaining with transfers, and obtain a completely Pareto unranked set of equilibria assuming that play under disagreement does not vary with the manner in which bargaining broke down. [Safronov and Strulovici \(2014\)](#) study a different protocol that does not satisfy such condition. Under this protocol, Pareto ranked equilibria can co-exist because players can be punished by the simple fact of making an offer, and, thus, profitable proposals might be deterred in equilibrium. Such punishments

⁷Communication among the firms might also play a role in the implementation phase after collusion has already been initiated. This might be the case when firms cannot monitor each other’s prices perfectly ([Rahman, 2014](#), [Awaya and Krishna, 2015](#)), or when they have private information about market conditions ([Athey and Bagwell, 2001](#)). However, none of these features are present in my model.

cannot occur in my model because, in the limit when revision opportunities are infrequent, the set of equilibrium continuation values after each specific proposal collapses to a singleton. Hence, the renegotiation protocol induced by the use of algorithms does guarantee that the set of equilibrium continuation values converges to a Pareto unranked set.

Asynchronous moves.— In my model, firms never revise their algorithms exactly at the same time. There are other papers in which asynchronous timing helps to reduce the set of equilibria. In particular, [Lagunoff and Matsui \(1997\)](#) obtain a unique equilibrium for *perfect coordination* asynchronous repeated games.⁸ More recently, [Calcagno et al. \(2014\)](#) consider a model with asynchronous revision opportunities, and establish uniqueness of equilibria for *single-shot* common interest games, and 2×2 opposing interest games. Regarding the study of oligopolies, [Maskin and Tirole \(1988a,b\)](#) study a class of asynchronous pricing games, and obtain unique equilibria restricting attention Markov strategies. [Eaton and Engers \(1990\)](#) study a different pricing game, but they focus on the possibility rather than the necessity of collusion. A salient aspect of my model is that firms do not choose simple acts, but rather algorithms that can react to market outcomes. This allows me to obtain sharp predictions for dynamic duopoly games without restricting attention to Markov strategies.

Choices via algorithms.— [Rubinstein \(1986\)](#) and [Abreu and Rubinstein \(1988\)](#) analyze repeated games in which players play via finite automata chosen at the beginning of time, and strictly prefer automata with fewer states. While, in some cases, these considerations narrow the set of equilibria, non-collusive equilibria remain. An important difference is that firms in my model are indifferent regarding the complexity of the algorithms they use. [Tennenholtz \(2004\)](#) studies a model in which players implement strategies via computer programs that can read each other codes. The set of Nash-equilibrium payoffs of this model coincides with the set of feasible and individually rational payoffs. Hence, in the context of oligopolies, such a programs could *enable* collusion, but need not guarantee it. Also, all the three mentioned papers focus on the case in which players choose programs or automata at the beginning of the game and never have a chance to revise them.

⁸A perfect coordination game is one in which players incentives are perfectly aligned, i.e., they share exactly the same preferences. [Yoon \(2001\)](#) shows that a folk theorem applies for any game that is not a perfect coordination game, even if choices are asynchronous.

2. Model

I consider a symmetric dynamic model of price competition with two firms $j \in \{1, 2\}$. Time is continuous and is indexed by $t \in [0, \infty)$. Consumers arrive randomly following a Poisson process with parameter $\lambda > 0$. Let $y = (y_n)_{n \in \mathbb{N}}$ denote the sequence of (random) consumer-arrival times. For simplicity, I assume that a single consumer arrives at each y_n .

2.1. Price competition

When a consumer arrives in the market, firms simultaneously offer prices p_1 and p_2 , respectively. The consumer observes both prices and decides whether to buy a *single unit* from one of the firms or to not buy at all. The consumer's decision depends on his type (ξ_1, ξ_2) , where ξ_j is the value of consuming the good produced by firm j . The consumer's type (ξ_1, ξ_2) is jointly normally distributed with mean zero, unit variance and correlation $\rho \in [0, 1)$. The utility from not buying is normalized to 0, and the net utility from buying from firm j is given by $\xi_j - p_j$. The "demand" for firm j is thus the probability that the consumer buys its product. Marginal costs are constant and normalized to 0, so that firm j 's expected profit is given by

$$\pi_j(p) = p_j \times \Pr(\xi_j - p_j > \max\{0, \xi_{-j} - p_{-j}\}). \quad (1)$$

Let \bar{p} be the price that firm j would charge optimally if it was the only firm in the market, i.e., it is the price that maximizes $p_j \times \Pr(\xi_j - p_j > 0)$. Let $\pi(p) = (\pi_1(p), \pi_2(p))$ and define $\Pi \subseteq \mathbb{R}_+^2$ by $\Pi = \{\pi(p) \mid p \in [0, \bar{p}]^2\}$. In words, Π would be the set of feasible profit profiles if firms could choose any prices on $[0, \bar{p}]$. My assumptions guarantee that Π is a convex and compact subset of \mathbb{R}_+^2 , and that its upper envelope is strictly concave. Consequently, for every point π^0 on the Pareto frontier of Π , there exists a price vector $p^0 \in [0, \bar{p}]^2$ such that $\pi^0 = \pi(p^0)$. Let p^M be the price that a monopolist who owns both firms would charge to maximize joint profits, and let $\pi^M = \pi(p^M, p^M)$. The minimax profit for each firm is $\underline{\pi} = \max_{p_j \in P} \pi_j(p_j, 0)$. The maximum profit for any firm is $\hat{\pi} = \max\{\pi_j \mid \pi \in \Pi\}$.

Prices are restricted to belong to the finite set

$$P = \left\{ 0, \frac{\bar{p}}{K_p}, \frac{2\bar{p}}{K_p}, \dots, \frac{(K_p - 1)\bar{p}}{K_p}, \bar{p} \right\},$$

where $K_p \in \mathbb{N}$ is an exogenous parameter. The set of feasible profits is thus a proper subset of Π , but its limit when K_p goes to infinity is dense in Π . The main results consider the limit as the number of points in the grid K_p goes to infinity.

Restricting attention to finite action spaces guarantees existence of equilibria but plays no role otherwise.

2.2. Algorithmic pricing

I assume perfect monitoring—each firm observes all past past consumer arrivals, prices and sales. Firms use *pricing algorithms* that automatically set prices as a function of the history of market outcomes. Pricing algorithms are modelled as finite automata characterized by tuples $a = (\Omega, \omega_0, \theta, \alpha) \in A$. Ω is a finite set of states, with $\|\Omega\| \leq K_a$ for some exogenous fixed bound $K_a \in \mathbb{N}$. $\omega_0 \in \Omega$ is the initial state. $\omega : \Omega \times P \times C \rightarrow \Omega$ is a transition function that depends on the current state, the price chosen by the competition, and the consumer’s decision. Finally, $\alpha(\omega) \in P$ is the price set by the algorithm every time a customer arrives and the current state is ω .




Figure 1 – Examples of pricing algorithms.

Figure 1 illustrates some examples of pricing algorithms. Each circle represents a state ω , and the price inside the circle corresponds to $\alpha(\omega)$. The initial state is denoted by a bold contour. The transitions are depicted by arrows. The algorithm “*always monopolistic*” always sets the price to p^M . The algorithm “*grim trigger*” sets the price p_j for as long as $-j$ offers p_{-j} , and switches to setting the price to 0 forever after the first time $-j$ sets any price other than p_{-j} . The algorithm “*two monopolistic*” sets the monopolistic price p^M for the first two consumers, and then offers the product for free to every subsequent consumer, regardless of market outcomes.

The transition functions for the algorithms depicted on Figure 1 only depend on the price set by the competition and the current state. In principle, the transitions could also depend on sales, but this would not affect the subsequent analysis.

2.3. Dynamic game

The dynamic game begins with firms choosing pricing algorithms simultaneously and independently at time 0. Although firms can perfectly observe each

other’s prices and sales, they cannot observe each other’s algorithms. However, over time, they either infer or are able successfully decode these. The time taken to decode is also random and independent across firms. So, for instance, after the initial choice of algorithms, firm 1 may be the first to decode firm 2’s algorithm at time x_{11} . At the time firm 1 decodes its rival’s current algorithm, it can choose to revise its own algorithm if it so wishes.

Following Kamada and Kandori (2009), I call such an event a *revision opportunity* at time x_{11} . Formally, revision opportunities for firm j arise stochastically according to a Poisson process with parameter $\mu_j > 0$. With the exception of Section 5.2, I only consider the symmetric case with $\mu_1 = \mu_2 =: \mu$. Not only is the arrival of revision opportunities independent across firms, it is also independent of the arrival of consumers. Let $x_j = (x_{jn})_{n \in \mathbb{N}}$ denote the sequence of signal-arrival times for firm j .

A feasible history of length N consists of a finite sequence of events—either consumer arrivals or revision opportunities—described by vectors of the form $h = (z_n, i_n, a_n)_{n=1}^N \in H_N$. $z_n \in [0, \infty)$ corresponds to the time of the n -th event, e.g., $z_1 = \min\{y_1, x_{11}, x_{21}\}$. $i_n \in \{0, 1, 2\}$ indicates the nature of each event, it takes the value 0 if the n -th event is a customer arrival, and the value $j = 1, 2$, if it is a revision opportunity for firm j and is set to 0 otherwise.⁹ Finally, $a_n \in A^2$ indicates the profile of algorithms being employed at the time of the n -th occurrence. Note that h_t is sufficient to derive the history of prices, profits and signals. Given a history $h \in H_N$, let $z = (z_n)_{n=1}^N$, $i = (i_n)_{n=1}^N$, $a = (a_n)_{n=1}^N$.

Each firm j must choose a pricing algorithm at time 0, and at each time it has a revision opportunity. Let H_{jN} denote the set of histories $h \in H_N$ such that $i_N = j$, and let $H_{j0} = \{\emptyset\}$ denote the initial history at time 0. The set of *decision nodes* for firm j is $H_j = \cup_{N=0}^{\infty} H_{jN}$. For tractability, I restrict attention to strategies that do not depend on the current time, nor on the specific arrival-times of signals and consumers. A (time-homogeneous behavior) *strategy* for firm j is a function $s_j : H_j \rightarrow \Delta(A)$ such that $s_j(h) = s_j(h')$ whenever $i = i'$ and $a = a'$, regardless of z and z' . Let S_j denote the set of strategies for firm j , and S the set of strategy profiles $s = (s_j, s_{-j})$. Continuation strategies are denoted by $s_j|_h \in S_j|_h$. The restriction to time-homogeneous strategies is discussed below.

⁹I am implicitly restricting attention to histories in which different events don’t occur at the same time. The set of such histories occurs with probability 1.

2.4. Expected discounted profits

Firms discount future profits with a common discount rate $r > 0$, and seek to maximize their expected discounted profits defined by

$$v(s) = (v_1(s), v_2(s)) := \frac{r}{\lambda} \times \mathbb{E}_s \left[\sum_{n=1}^{\infty} \exp(-ry_n) \pi_j(p_n) \right],$$

where p_n is the price vector offered to the n -th consumer to arrive. The term r/λ is a normalizing factor that guarantees that expected discounted profits are expressed in the same units as the stage-game profits, so that $v(s) \in \Pi$. I employ two convenient expressions for v , that are possible because of the time-homogeneity assumption and the properties of independent Poisson processes. The derivation of such expressions is in Appendix A.1.

First, average expected discounted profits can be expressed as

$$\begin{aligned} v(s) &= \frac{\lambda}{\lambda + 2\mu + r} \pi^1 + \frac{\lambda}{\lambda + 2\mu + r} w^0 + \frac{\mu}{\lambda + 2\mu + r} w^1 + \frac{\mu}{\lambda + 2\mu + r} w^2 \\ &= (1 - \delta) \pi + \delta \left[\frac{\lambda}{\lambda + 2\mu} w^0 + \frac{\mu}{\lambda + 2\mu} w^1 + \frac{\mu}{\lambda + 2\mu} w^2 \right]. \end{aligned} \quad (2)$$

where π^1 is the vector of expected profits from the first consumer to arrive given the initial pricing algorithms, w^i is the vector of continuation values after the first event for $i \in \{0, 1, 2\}$, and δ is the *effective discount factor* between consumer arrivals defined as

$$\delta := \mathbb{E}[\exp(-ry_1)] = \frac{\lambda}{\lambda + r}. \quad (3)$$

Now, suppose that the current profile of firms' pricing algorithm is a , and let s_j be any strategy for firm j that chooses a_j in any history before firm $-j$'s first revision. In that case, the average expected discounted profits of the firms can be decomposed as

$$v(a, s) = \frac{r}{\mu + r} \pi(a; \beta) + \frac{\mu}{\mu + r} \tilde{w}, \quad (4)$$

where $\pi(a; \beta)$ corresponds the expected discounted profits that would result if no new revision opportunities ever arrived, and \tilde{w} is the vector of expected discounted continuation values after $-j$'s first revision. Formally, $\pi(a; \beta)$ and \tilde{w} are defined as

$$\pi(a; \beta) := (1 - \beta) \sum_{n=0}^{\infty} \beta^n \pi(p_{n+1}(a)) \quad \text{and} \quad \tilde{w} := (1 - \beta) \sum_{n=0}^{\infty} \beta^n w_n, \quad (5)$$

where $p_n(a)$ is the vector of prices offered to the n -th consumer to arrive according to a , w_n is the vector of continuation values at the moment of $-j$'s next revision if it comes exactly after the n -th consumer, and

$$\beta := \frac{\lambda}{\lambda + \mu + r}. \quad (6)$$

Not that, in both (2) and (4), v is expressed as a convex combination of different terms, and the corresponding weights are a function the ratios μ/r and r/λ . The results in the rest of the paper refer to limiting cases when these ratios are close to 0 or diverge to ∞ , and they obtain because, in such cases, some of the terms in these expressions dominate the others. For example, in Section 4, I consider the limit when $\mu/r \rightarrow 0$ and $r/\lambda \rightarrow 0$. In that case, v is approximately equal to w_0 in (2), and approximately equal to $\pi(a; \beta)$ in (4). This implies that, in the limit, firms only care about playing a best response to the current algorithm of their opponents.

Restricting attention to time-homogeneous strategies allows me to use the simple expressions from (2) and (4). The analogous expressions with general strategy spaces would be significantly less tractable, but one should expect that similar arguments should work. For example, in the limit when $\mu/r \rightarrow 0$ and $r/\lambda \rightarrow 0$, firms would still only care about playing a best response to the current algorithm of their opponents. However, the derivations would become significantly more cumbersome.

2.5. Solution concept

Recall that there is perfect monitoring of market outcomes, and perfect observability of algorithms at every revision opportunity. Thus, the dynamic game is a game of complete perfect information, except for the first period when initial algorithms are chosen simultaneously. Consequently, it suffices to consider subgame perfect equilibria in time-homogeneous strategies. Let S^* denote the set of equilibria and $V^* = \{v(s) \mid s \in S^*\}$ the set of feasible equilibrium profits.

The example from Section 3 ahead shows that unconditional repetition of a static Nash equilibrium of the stage game might not constitute an equilibrium of the dynamic game. Hence, establishing existence of equilibria is not completely straightforward. However, the restrictions on the restrict the complexity of pricing algorithms and the set of feasible prices guarantee that the dynamic game is a stochastic game with a *finite* state space, and hence an equilibrium exists. A direct proof is provided in Appendix A.2.

Proposition 2.1 *The dynamic game always admits a symmetric equilibrium.*

3. A two-price example

The following example illustrates the mechanism through with renegotiation helps to eliminate inefficient equilibria. Suppose that the set of prices is restricted to be $P = \{p^B, p^M\}$ and expected profits π are as in Figure (2). These profits could arise, for instance, if p^B is the (Bertrand) equilibrium price of the unconstrained game with $P = [0, \bar{p}]$. Note that the stage game corresponds to a prisoner's dilemma with p^B being strictly dominant.

	p^M	p^B
p^M	3, 3	0, 4
p^B	4, 0	1, 1

Figure 2 – Profit matrix for two-price example.

Now suppose that $K_a \geq 2$, $\delta > 1/3$ and μ is small enough so that $\beta > 1/3$ and

$$\mu < \frac{1}{4}\beta^n(1 - \beta)(3\beta - 1)r \quad \text{and} \quad \mu < (3(1 - \beta^n) - 1)r,$$

for some number $n > \log(2/3)/\log \beta$.¹⁰ I will show that, in this case, there does *not* exist an equilibrium in which the price along the equilibrium path is always p^B . Consequently, under these parameter conditions, the firms' joint profits in equilibrium need be strictly greater than the competitive profits. Suppose towards a contradiction that there exists one such equilibrium, say, s^B . Since $K_a \geq 2$, firms 1 could choose at time 0 the two-state algorithm “*tit for tat*” depicted in Figure 3. Note that any potential differences in profits between this choice and s^B need to occur after firm 2's first revision opportunity.




Figure 3 – Pricing algorithms for the two-price example.

Now, let us think of what happens when firm 2 has a revision opportunity and

¹⁰To see that there exist such μ and n , fix any $\delta < 1/3$ and set $n = 1 + \log(2/3)/\log \delta$. Note that $\lim_{\mu \rightarrow 0} \beta = \delta$. Hence, the bound for n is satisfied for μ small enough, and guarantees that the right-hand sides of the bounds for μ have strictly positive limits.

observes that firm 1 is using “*tit for tat*”. If firm 2 chooses “*always monopolist*”, we know from (4) that its continuation value would be bounded below by

$$v_2' = \frac{r}{\mu + r}((1 - \beta)0 + \beta 3) + \frac{\mu}{\mu + r}0 = \frac{3\beta r}{\mu + r}.$$

Suppose in contrast that firm 2 chooses a continuation automaton a_2^n that doesn't choose p^M for at least the first n consumers. Using (4) and the fact that $\beta > 1/3$, firm 2's continuation value would be bounded above by

$$\begin{aligned} \hat{v}_2^n &= \frac{r}{\mu + r}\pi_2(\textit{tit for tat}, a_2^n) + \frac{\mu}{\mu + r}\tilde{w} \\ &\leq \frac{r}{\mu + r}\left[3(1 - \beta^{n-1}) + \beta^n(1 - \beta)4 + \beta^{n+2}3\right] + \frac{\mu}{\mu + r}4. \end{aligned}$$

The condition $\mu < (1 - \beta)(3\beta - 1)r/4$ implies, after some simple algebra, that $v_2' > \hat{v}_2^n$. Hence, in any equilibrium, after observing “*tit for tat*”, firm 2 must choose an algorithm that sets the monopolistic price for at least n periods.

Therefore, using (4) once again, it follows that, in any equilibrium, firm 1's continuation value after such history satisfies

$$v_1 \geq \frac{r}{\mu + r}\left[(1 - \beta)4 + \beta(1 - \beta^{n-1})3\right] > 1,$$

where the last inequality follows from $\mu > [3(1 - \beta^n) - 1]r$. This implies that, by choosing “*tit for tat*” at time $t = 1$, firm 1 can guarantee a payoff strictly greater than 1. Since $v_1(s^B) = 1$, it follows that s^B cannot be an equilibrium.

The firm's ability to commit and to decode their rival's algorithms is not sufficient to preclude low-profit equilibria. Suppose that $K_a = 1$, so that the only available algorithms are “*always monopolistic*” and “*always Bertrand*”. In this case, the strategy profile according to which both firms choose “*always Bertrand*” after any history constitutes an equilibrium of the dynamic game. This is because, if future play does not depend on the firms' current choices, they have no reason to deviate from their dominant price. Hence, it is crucial that firms can choose sophisticated algorithms that react to market conditions. Now suppose that $\mu = 0$, so that firms choose algorithms at the beginning of the game and cannot revise them ever after. If firm 1 expects firm 2 to choose “*always Bertrand*”, then choosing “*always Bertrand*” is a best response. Hence, it is also crucial that firms can revise their algorithms over time (and thus commitment is imperfect).

4. Inevitability of collusion

The previous example illustrates how pricing algorithms can provide a renegotiation channel that precludes low-profit equilibria. Now, I will show that in some limiting parameter regions when consumers arrive frequently and revision opportunities are infrequent, the power of renegotiation is such that *all* equilibria of the game lead to joint continuation profits that are *arbitrarily* close to monopolistic profits. Formally, the main result reads as follows.

Theorem 4.1 [Inevitability of collusion] *Fix any r . For any $\nu > 0$ the probability that, in any equilibrium, the joint expected discounted continuation profits are farther away than ν from the joint monopolistic profits converges to zero in the limit when revision opportunities are arbitrarily infrequent and costumers arrive arbitrarily frequently, i.e.,*

$$\lim_{\lambda \rightarrow \infty} \limsup_{\mu \rightarrow 0} \sup_{s^* \in S^*} \Pr_{s^*} \left(\bar{\pi}^M - \bar{v}(s^*|_h, h) > \nu \mid h \in H^2 \right) = 0,$$

where H^2 is the set of histories in which each firm has had at least one revision opportunity.

Establishing Theorem 4.1 requires a number of steps. First, I show that, in the limit when revisions are infrequent and consumers arrive frequently, some collusive proposals need to be accepted. Formally, lemmas 4.2 and 4.3 establish a lower bound for the continuation values after some histories in which a firm observes that its rival is using a specific kind of algorithms specified below. Then, I show that, after having a revision opportunity, a firm can include such proposals into its algorithm to guarantee high continuation values. As a consequence, the firms continuation values after subsequent revisions need to be close to Pareto frontier (Lemma 4.4). Finally, I show that such continuation values are *sufficiently* close to the Pareto frontier to guarantee that joint profits will remain high with probability approaching one. The lemmas are formally stated and explained with more detail below. The formal proof of the lemmas and the main result are in Appendix B.

4.1. Main steps of the proof

The constructions below are based on grim trigger algorithms corresponding to strictly individually rational profit profiles along the Pareto frontier. Formally, given a price vector $p^0 \in P$, the corresponding *grim trigger algorithm* for firm j is the algorithm that plays p_j^0 as long as it sees p_{-j}^0 and plays 0 forever after any other history, as depicted in Figure 1. Say that p^0 is *strictly individually rational*

if $\Delta(p^0, \beta) > 0$, where

$$\Delta(p^0, x) := \min_{j=1,2} \left\{ \pi_j(p^0) - [(1-x)\hat{\pi}_j(p_0) + x\bar{\pi}] \right\}.$$

This is equivalent to say that the corresponding grim-trigger algorithm profile would be a strict equilibrium of a repeated game with discount factor β . A sufficient condition for p^0 to be strictly individually rational is that $\pi_j(p^0) > \bar{\pi}$ for $j = 1, 2$, and β is sufficiently close to 1.

When firm j observes that its rival is using the grim trigger algorithm corresponding to p^0 , it has two alternatives. It can choose an algorithm that plays p_j^0 for a “long period of time”, thus guaranteeing the profit of $\pi_j(p^0)$ until $-j$'s next revision. Otherwise, it can choose a different algorithm that might yield a high profit the first time it deviates from p_j^0 , but, after that, it would trigger $-j$'s grim trigger and result in profits of at most $\bar{\pi}$ until $-j$'s next revision. If p^0 is strictly individually rational and μ is close to 0, it is very unlikely that $-j$'s next revision happens soon and, therefore, the first option is preferable.

Exactly how long is this “long period of time”? The following lemma states that firm j will choose an algorithm that offers p_j^0 to a number of consumers which goes to ∞ at the rate of $1/\mu$ when μ goes to 0. In order to establish this bound, first I establish a weaker bound, based on equation (4), which also diverges but at a slower rate. Then, I use this bound to show that the gap between different possible continuation values after $-j$'s revisions is small. This implies that firm j is even more concerned about playing a best response to $-j$'s current algorithm. Hence, firm j will play p_j^0 for even longer, which, in turn, implies that the gap in continuation values is even smaller, and so on. The bound from Lemma 4.2 is the limit that obtains from iterating this process indefinitely.

Lemma 4.2 [An offer that cannot be refused] *Fix any r and λ , and a strictly individually rational price vector p^0 such that $\pi(p^0)$ belongs to the Pareto frontier of the set of feasible profits, and let a^0 be the corresponding grim-trigger algorithm profile. In any equilibrium, after firm j has a revision opportunity and observes a_{-j}^0 , it choose an automaton that offers p_j^0 as long as it continues to observe p_{-j}^0 for at least n periods, where n is the largest integer satisfying*

$$n \leq N(\pi^0; \lambda, \mu, r) := \frac{r\Delta(p^0, \beta)}{2\mu\hat{\pi}} - \frac{\mu}{2(\mu + r)}. \quad (7)$$

The previous Lemma guarantees that, after observing a grim trigger algorithm corresponding to a Pareto efficient, strictly individually rational price vector p^0 , firm j will mimic a dynamic best response for a long period of time. This results in a lower bound for firm $-j$'s continuation profit after such history, that approx-

imates $\pi_{-j}(p^0)$. Remarkably, the approximation is extremely good. The lower bound from the following lemma converges to $\pi_{-j}(p^0)$ at the rate of $\exp(-1/\mu)$ when μ goes to 0.

Lemma 4.3 [Approximation rate] *Let p^0 , π^0 and a^0 be as in Lemma 4.2. For any $r > 0$, there exists some $\underline{\mu}(r) > 0$, which does not depend on p^0 , such that, if $\mu \leq \underline{\mu}(r)$ and $\Delta(p^0, \beta) > 0$, then, in any equilibrium, firm j 's continuation value after a history in which firm $-j$ observes a_j^0 can be no worse than $\pi_j^0 - \epsilon(\lambda, \mu, r)$, where*

$$\epsilon(\lambda, \mu, r) := \frac{1}{\delta} \hat{\pi} \exp\left(-C_1(r) \frac{\Delta(p^0, \beta)}{\mu}\right), \quad (8)$$

where $C_1(r) > 0$ is a positive number that does not depend on p^0 or μ .

Now consider a history starting from a revision by firm j . Consider a continuation strategy profile in which the continuation values after $-j$'s next revision are Pareto dominated, with high probability, by some individually rational and Pareto efficient profits $\pi^0 = \pi(p^0)$. Firm j could deviate by appending its automaton with a grim trigger algorithm for p^0 , to be activated the first time $-j$ uses the price p_{-j}^0 . Informally, this can be thought of as using an algorithm that makes a collusive offer. The previous two lemmas can be used to guarantee that $-j$ would accept this offer for sufficiently small μ , and hence the deviation would be profitable. Formalizing this argument results in the following lemma. The lemma asserts that, at the time of each subsequent revision after both firms have had at least one revision opportunity, continuation values need to be close to the Pareto frontier with high probability.

Lemma 4.4 [Efficient renegotiation] *Fix any $r > 0$ and $\mu \leq \underline{\mu}(r)$ where $\underline{\mu}(r)$ is the bounds from Lemma 4.3. Consider any equilibrium and a history at which firm j has a revision opportunity, and let \bar{w}' denote firms joint expected discounted continuation profits at the time of of firm $-j$'s next revision. For any $m > 0$ the probability that the difference between these joint continuation profits and joint monopolistic profits is greater than $m\epsilon(\lambda, \mu, r)$ is bounded above by*

$$\Pr\left(\bar{\pi}^M - \bar{w}' > m\epsilon(\lambda, \mu, r)\right) \leq \frac{1}{m+1}. \quad (9)$$

An additional step is needed to establish Theorem 4.1. The previous lemma guarantees that continuation profits will be high at the moment of subsequent revisions. However, there is a possibility that firms use algorithms that use high prices right after each revision, but lower and lower prices as times goes on. Hence,

continuation profits could gradually decrease between revisions. Fortunately, the rate from Lemma 4.3 is extremely fast. This guarantees that, for any constant ν , the probability that joint continuation profits reach a point farther than ν from the joint monopolistic profits converges to 0 in the limit when μ goes to 0.

5. Additional results

5.1. Asymmetry and leadership

Proposition 5.1 *For every equilibrium of the dynamic game, there is an equilibrium of the alternative game which yield the same path of play and in which firms always choose to make their algorithms transparent.*

5.2. Asymmetry and leadership

In this section, I consider the asymmetric case when firm 1 is very committed ($\mu_1 \rightarrow 0$), while firm 2 revises frequently ($\mu_2 \rightarrow \infty$). In this case, firm 1 acts like a Stackelberg leader that makes an ultimatum offer at the beginning of the game, and thus extracts all the market profits. Formally, I will show that, in the mentioned limit, firm 1's worst equilibrium value approximates its dynamic Stackelberg payoff defined ahead.

For any automaton a let $(p_n(a))$ be the sequence of induced prices, and let $\pi(a; \delta)$ be the expected discounted profits if no revisions arise, defined as in (5). Say that a_{-j} is a best response to a_j in the game with no revisions, if $\pi_j(a; \delta) > \pi_j(a'_j, a_{-j}; \delta)$ for all $a'_j \in A$. Define the *dynamic Stackelberg profit* for form i as follows

$$\pi_i^S(\lambda, r) = \sup \{ \pi_j(a; \delta) \mid a \in A \text{ and } a_{-j} \text{ is a best response to } a_j \}.$$

In words, $\pi_i^S(\lambda, r)$ is the best payoff that firm j could receive if it committed to use an algorithm forever at date $t = 0$, and firm $-j$ chose its algorithm after observing this commitment. The proof of the proposition is in Appendix C.1.

Proposition 5.2 *In the limit when firm 1 is completely committed and firm 2 can revise arbitrarily often, firm 1's expected discounted profits in any equilibrium are weakly greater than its dynamic Stackelberg payoff, i.e., for any λ and r ,*

$$\lim_{\mu_1 \rightarrow 0} \lim_{\mu_2 \rightarrow \infty} \inf_{s^* \in S^*} v_1(s^*) \geq \pi_1^S(\lambda, r).$$

As usual, $\pi_i^S(\lambda, r)$ is weakly greater than the Bertrand profit from the stage

game. Moreover, as long as $K_A \geq 2$, it converges to

$$\pi^* := \sup\{\pi_j \mid \pi \in \Pi \text{ and } \pi_{-j} \geq \underline{\pi}\}$$

when δ goes to 1. Hence, Proposition 5.2 implies as a corollary that there is a *unique* equilibrium value in the limit when firms are very patient and asymmetrically committed. Formally,

$$\lim_{\lambda \rightarrow \infty} \lim_{\mu_1 \rightarrow 0} \lim_{\mu_2 \rightarrow \infty} V^*(\lambda, r, \mu_1, \mu_2) = \left\{ \begin{pmatrix} \pi^* \\ \underline{\pi} \end{pmatrix} \right\}.$$

5.3. Collusion among impatient firms

The rest of the paper focuses on the ability to renegotiate when firms are sufficiently inflexible ($\mu \rightarrow 0$). In contrast, this section consider the limiting case when revision opportunities are arbitrarily frequent ($\mu \rightarrow \infty$). In this case, the use of *observable* pricing algorithms can foster collusion through a different mechanism, by improving the firms' ability to monitor each other. The improved monitoring facilitates the enforcement of collusive agreements, even if firms are extremely impatient.

Proposition 5.3 *Firms can guarantee profit arbitrarily close to the symmetric efficient monopolistic outcome when revision opportunities are sufficiently frequent relative to consumer-arrivals, i.e., for any r and λ ,*

$$\lim_{\mu \rightarrow \infty} \sup \{v \mid (v, v) \in V^*(\lambda, \mu, r)\} = \pi^M.$$

The salient feature of Proposition 5.3 is that collusion is possible *for any* discount factor δ . In contrast, in the repeated game (with $\mu = 0$), sustaining monopolistic profits in equilibrium is only possible if δ is high enough so that $\Delta(p^M, \delta) \geq 0$. Otherwise, an agreement to sustain π^M would break down because firms would prefer to deviate in order to make a positive profit in the short run. These deviations are no longer possible when revision opportunities are frequent enough, because, in this case, there is a high probability that any potential deviation will be detected even before the next customer arrives. Hence, the revising firm could react to the deviation before the deviating firm has an opportunity to make any profits. The proof of Proposition 5.3 is in appendix C.2.

6. Conclusion

My findings provide theoretical support for the idea that optimal use of pricing algorithms is an effective tool for tacit collusion. When firms set prices through algorithms that can respond to market conditions, are fixed in the short run, can be decoded by rivals, and can be revised over time, every equilibrium of the game leads in the long run to monopolistic profits. It is this the combination of all these features that makes collusion inevitable, neither ingredient on its own will yield the result. While the analysis is carried out in the context of a pricing duopoly, it is clear that the lessons are more general and should apply to the context of general repeated games.

References

- Abreu, D. and Rubinstein, A. (1988). The structure of nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6):1259–1281.
- Anderson, R. M. (1984). Quick-response equilibrium. Paper IP 323, Center for Research in Management, University of California Berkeley.
- Athey, S. and Bagwell, K. (2001). Optimal collusion with private information. *The RAND Journal of Economics*, 32(3):428–465.
- Aumann, R. J. and Sorin, S. (1989). Cooperation and bounded recall. *Games and Economic Behavior*, 1(1):5–39.
- Awaya, Y. and Krishna, V. (2015). On communication and collusion. *The American Economic Review*. Forthcoming.
- Bernheim, B. D. and Ray, D. (1989). Collective dynamic consistency in repeated games. *Games and Economic Behavior*, 1(4):295–326.
- Block, J. I. and Levine, D. K. (2015). Codes of conduct, private information, and repeated games. *International Journal of Game Theory*. Forthcoming.
- Boehmer, E., Li, D., and Saar, G. (2015). Correlated high-frequency trading. Mimeo.
- Borenstein, S. (2004). Rapid price communication and coordination: The airline tariff publishing case (1994). In Kwoka, Jr., J. E. and Finnegan, N. F., editors, *The Antitrust Revolution: Economics, Competition, and Policy*, pages 233–251. Oxford University Press.
- Calcagno, R., Kamada, Y., Lovo, S., and Sugaya, T. (2014). Asynchronicity and coordination in common and opposing interest games. *Theoretical Economics*, 9(2):409–434.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.
- Eaton, J. and Engers, M. (1990). Intertemporal price competition. *Econometrica*, 58(3):637–659.
- Ezrachi, A. and Stucke, M. E. (2015). Artificial intelligence & collusion: When computers inhibit competition. Oxford Legal Studies Research Paper 18/2015.
- Farrell, J. and Maskin, E. (1989). Renegotiation in repeated games. *Games and Economic*

- Behavior*, 1(4):327–360.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12.
- Green, E. J., Marshall, R. C., and Marx, L. M. (2014). Tacit collusion in oligopoly. In Blair, R. D. and Sokol, D. D., editors, *The Oxford Handbook of International Antitrust Economics*, volume 2, pages 464–497. Oxford University Press.
- Harrington, J. E. (2015). A theory of collusion with partial mutual understanding. Mimeo.
- Ivaldi, M., Jullien, B., Rey, P., Seabright, P., and Tirole, J. (2003). The economics of tacit collusion. Final report for DG competition, European Commission.
- Kamada, Y. and Kandori, M. (2009). Revision games. Mimeo.
- Kamada, Y. and Kandori, M. (2011). Asynchronous revision games. Mimeo.
- Lagunoff, R. and Matsui, A. (1997). Asynchronous choice in repeated coordination games. *Econometrica*, 65(6):1467–1477.
- Levine, D. K. and Pesendorfer, W. (2007). The evolution of cooperation through imitation. *Games and Economic Behavior*, 58(2):293–315.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press.
- Maskin, E. and Tirole, J. (1988a). A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica*, 56(2):549–569.
- Maskin, E. and Tirole, J. (1988b). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56(2):571–599.
- Mehra, S. K. (2015). Antitrust and the robo-seller: Competition in the time of algorithms. *Minnesota Law Review*, 100. Forthcoming.
- Miller, D. A. and Watson, J. (2013). A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350.
- Pearce, D. (1987). Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation. Discussion Paper 855, Yale Cowles Foundation.
- Rahman, D. (2014). The power of communication. *The American Economic Review*, 104(11):3737–3751.
- Rubinstein, A. (1986). Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 39(1):83–96.
- Safronov, M. and Strulovici, B. (2014). Explicit renegotiation in repeated games. Mimeo.
- Tennenholtz, M. (2004). Program equilibrium. *Games and Economic Behavior*, 49(2):363–373.
- Yoon, K. (2001). A folk theorem for asynchronously repeated games. *Econometrica*, 69(1):191–200.

A. Omitted details from Section 2

A.1. Derivation of expected discounted profit representations

The fact that revision opportunities and customer arrivals are independent Poisson processes guarantees that they satisfy two desirable properties. First, z is itself a Poisson process with parameter $(\lambda + 2\mu)$. Second, i is a sequence of i.i.d. variables independent

of z and distributed according to

$$\Pr(i_n = 0) = \frac{\lambda}{\lambda + 2\mu} \quad \text{and} \quad \Pr(i_n = j) = \frac{\mu}{\lambda + 2\mu} \quad \text{for } j = 1, 2.$$

Moreover, time-homogeneity implies that prices and continuation values are also independent of z_t . Hence, we can derive (2) as follows

$$\begin{aligned} v &= \mathbb{E} \left[\exp(-rz_1) \left(\mathbf{1}(i_1 = 0) \left(\frac{r}{\lambda} \pi(p_1) + w^0 \right) + \mathbf{1}(i_1 = 1)w^1 + \mathbf{1}(i_1 = 2)w^2 \right) \right] \\ &= \mathbb{E}[\exp(-rz_1)] \left(\Pr(i_1 = 0) \left(\frac{r}{\lambda} \pi_1 + w^0 \right) + \Pr(i_1 = 1)w^1 + \Pr(i_1 = 2)w^2 \right) \\ &= \frac{\lambda + 2\mu}{\lambda + 2\mu + r} \left(\frac{\lambda}{\lambda + 2\mu} \left(\frac{r}{\lambda} \pi_1 + w^0 \right) + \frac{\mu}{\lambda + 2\mu} w^1 + \frac{\mu}{\lambda + 2\mu} w^2 \right) \\ &= \frac{r}{\lambda + 2\mu} \pi_1 + \frac{\lambda}{\lambda + 2\mu} w^0 + \frac{\mu}{\lambda + 2\mu} w^1 + \frac{\mu}{\lambda + 2\mu} w^2. \end{aligned}$$

In order to derive (4), let w_n^0 denote the vector of continuation values after n consumers have arrived if $-j$ has had no revision opportunities, in particular $v = w_0^0$. Since firm j will not revise its automaton until $-j$ does, it follows that $w^j = v$. Hence, using (2), we can write

$$v = \frac{r}{\lambda + \mu + r} \pi(p_1(a)) + \frac{\lambda}{\lambda + \mu + r} w_1^0 + \frac{\mu}{\lambda + \mu + r} w_0^{-j}.$$

By a similar argument, for any $n \geq 0$ we can write

$$w_n^0 = \frac{r}{\lambda + \mu + r} \pi(p_{n+1}(a)) + \frac{\lambda}{\lambda + \mu + r} w_{n+1}^0 + \frac{\mu}{\lambda + \mu + r} w_n^{-j}. \quad (10)$$

Equation (4) follows from iteratively applying (10), which yields

$$\begin{aligned} v &= \frac{r}{\lambda + \mu + r} \sum_{k=0}^{\infty} \left(\frac{\lambda}{\lambda + \mu + r} \right)^k \pi(p_{k+1}(a)) + \frac{\mu}{\lambda + \mu + r} \sum_{k=0}^{\infty} \left(\frac{\lambda}{\lambda + \mu + r} \right)^k w_k^{-j} \\ &= \frac{r}{\mu + r} (1 - \beta) \sum_{k=0}^{\infty} \beta^k \pi(p_{k+1}(a)) + \frac{\mu}{\mu + r} (1 - \beta) \sum_{k=0}^{\infty} \beta^k w_k^{-j} \\ &= \frac{r}{\mu + r} \pi(a; \beta) + \frac{\mu}{\mu + r} \tilde{w}. \end{aligned}$$

A.2. Existence of equilibria

Proof of Proposition 2.1. A strategy is *stationary* if $s_j(h_N) = s_j(h'_N)$ whenever $a_{-jN} = a'_{-jN}$, that is, the choice of firm j when it has a revision opportunity only depend on the algorithm currently being employed by its competitor. I will prove that an equilibrium exists within the class of stationary strategies. A stationary strategy s_j can be described by a an initial distribution $\sigma_{j0} \in \Delta(A_j)$ together with a transition function $\sigma_j : A_{-j} \rightarrow \Delta(A_j)$. Let $V(a|s)$ denote the expected continuation value when the current state is a and future choices are made according to s . These values are given by the unique solution

to

$$V(a|s) = \frac{\delta\mu_0}{\bar{\mu}} \left(\pi(a) + V(a^+|s) \right) \\ + \frac{\delta\mu_1}{\bar{\mu}} \sum_{a'_1} \sigma_2(a'_1) V(a'_1, a_2|s) + \frac{\delta\mu_2}{\bar{\mu}} \sum_{a'_2} \sigma_1(a'_2) V(a_2, a'_2|s).$$

And the expected discounted profits at time 0 are given by

$$V(s) = \sum_a \sigma_0(a) V(a).$$

To prove existence we define an auxiliary normal form game with players $I = J \times (\{\emptyset\} \cup A)$, strategy spaces $Z_i = A$, and utility function g described ahead. Given a strategy profile for the auxiliary game z , let s^z be the stationary strategy with $\sigma_{j0} = z_{(j0)}$ and $\sigma_j(a_{-j}) = z_{(j, a_{-j})}$, and let g_i with $i = (j, a_{-j})$ be given by

$$g_i(z) = V_j(s^z(a_{-j}), a_{-j}|s^z)$$

Because the auxiliary game is finite, it has a Nash equilibrium z^* (possibly in mixed strategies, redefine things properly).

Now let $s^* = s^{z^*}$. I claim that s^* is an equilibrium. Suppose not, then by the single deviation principle there is a single profitable deviation. The thing to prove is that this implies that there is a uniform deviation. By construction, a player cannot make a *uniform* deviation so we would be done. ■

B. Proof of the main result

The current version of the proof is missing some steps that will be added on subsequent drafts. First, the proofs are currently written for the limiting case when utility shocks become perfectly correlated ($\rho \rightarrow 1$) and thus goods are perfect substitutes. Second, since the convergence rate from Lemma 4.3 is not uniform, an additional step is needed to bound the difference between the best continuation value that is possible in equilibrium and the best continuation value that can be guaranteed in equilibrium.

Proof of Lemma 4.2. The following discussion applies at any history in which firm has a revision opportunity and observes the offer π^0 . Let N_j denote the minimum amount of periods for which the offer has to be accepted, let \underline{w}_j be the *worse* possible equilibrium continuation value starting from $-j$'s revision when j is using a_j^0 , and let \hat{w}_j be the *best* possible expected discounted continuation value for j after $-j$'s next revision. First, I will establish a lower bound for N_j that depends on the difference $(\hat{w}_j - \underline{w}_j)$. Then, I will establish an upper bound for this difference that depends on N_j and N_{-j} . Condition (7) results from combining these bounds.

For the first step, I will first finding a lower bound for j 's profits if it chooses a_j^0 and compare it with an upper bound if it decides to deviate from a_j^0 after a certain number of consumer have arrived. For the former bound note that, if $-j$ unconditionally sticks

to a_j^0 forever, it guarantees an expected discounted profit v_j satisfies

$$v_j \geq \underline{v}_j^{(\infty)} := \frac{r}{\mu + r} \pi_j^0 + \frac{\mu}{\mu + r} \underline{w}_j. \quad (11)$$

Now, let $\hat{v}_j^{(n)}$ denote the *best* possible continuation value for j starting from a history in which j is using a_j^0 , and j plans to stick to p_j^0 for the following n periods as long as $-j$ doesn't revise and deviate afterwards. Using (4), we can write

$$\hat{v}_j^{(n)} \leq \frac{r}{\mu + r} \left[\pi_j^0 - \beta^n \Delta_j(p^0, \beta) \right] + \frac{\mu}{\mu + r} \hat{w}_j. \quad (12)$$

In equilibrium, j must choose an automaton that plays p_j^0 for at least n periods whenever $\underline{v}_j^{(\infty)} \geq \hat{v}_j^{(n)}$. Using (11) and (12), this condition can be expressed as

$$\frac{r}{\mu + r} \pi_j^0 + \frac{\mu}{\mu + r} \underline{w}_j \geq \frac{r}{\mu + r} \left[\pi_j^0 - \beta^n \Delta_j(p^0, \beta) \right] + \frac{\mu}{\mu + r} \hat{w}_j.$$

which yields after some algebra

$$n \geq \frac{1}{\log \beta} \left[\log \left(\frac{\mu}{r} \right) + \log (\hat{w}_j - \underline{w}_j) - \log \Delta_j(p^0, \beta) \right]. \quad (13)$$

Hence, we obtain the following bound for N_j

$$\beta^{N_j} \leq \frac{\mu}{r \Delta_j(\beta)} (\hat{w}_j - \underline{w}_j) \quad (14)$$

In order to bound \hat{w}_j , note that continuation values satisfy

$$\hat{w}_j \leq \pi_j^0 + \kappa (\pi_{-j}^0 - \underline{w}_{-j}) + g(\pi_{-j}^0 - \underline{w}_{-j}). \quad (15)$$

where $\kappa = (\hat{\pi}_j^* - \pi_j^0) / (\pi_{-j}^0 - \underline{\pi}^*)$, and g satisfies, $g(0) = g(\pi_{-j}^0 - \underline{\pi}_j) = 0$, $\|g\|_\infty = \bar{g} < \infty$ and $g'' < 0$ for $\rho < 1$. In the limit case $\rho \rightarrow 1$, we have $g \equiv 0$.

Suppose that $-j$ has a revision opportunity and j 's current algorithm mimics a_j^0 for m periods with $0 \leq m \leq N_{-j}$. By choosing a_{-j}^0 , firm $-j$ guarantees a continuation profit of

$$\begin{aligned} \underline{w}_{-j}^{(m)} &\geq \frac{r}{\mu + r} (1 - \beta^m) \pi_{-j}^0 + \frac{\mu}{\mu + r} (1 - \beta^{N_j}) \pi_{-j}^0 \\ &= \pi_{-j}^0 - \left(\frac{r}{\mu + r} \beta^m + \frac{\mu}{\mu + r} \beta^{N_j} \right) \pi_{-j}^0 \end{aligned} \quad (16)$$

As before, the first term of the right-hand side corresponds to the profits that $-j$ would secure if j never had a new revision opportunity. The second term corresponds to the profits after j has a revision. The inequality follows because, after a revision, j will mimic a_j^0 for at least N_j customers.

For now let me focus on the limiting case in which $g \equiv 0$. From (15) and (16), it

follows that, for $k \leq N_j$ we have that

$$\hat{w}_j^{(N_j-k)} \leq \pi_j^0 + \kappa \left(\frac{r}{\mu+r} \beta^{-k} + \frac{\mu}{\mu+r} \right) \beta^{N_j} \pi_{-j}^0.$$

Also we know from construction that

$$\hat{\pi}_j^* = \kappa(\pi_{-j}^0 - \underline{\pi}^*) + \pi_j^0.$$

Hence, it follows that

$$\begin{aligned} \hat{w}_j &= (1-\beta) \sum_{k=0}^{N_j} \beta^k \hat{w}_j^{(n-k)} + (1-\beta) \sum_{k=N_j}^{\infty} \beta^k \hat{\pi}_j^* \\ &\leq \pi_j^0 + \kappa(1-\beta) \sum_{k=0}^{N_j} \beta^k \left(\frac{r}{\mu+r} \beta^{-k} + \frac{\mu}{\mu+r} \right) \beta^{N_j} \pi_{-j}^0 + \kappa(1-\beta) \sum_{k=N_j}^{\infty} \beta^k \pi_{-j}^0 \\ &= \pi_j^0 + \kappa \beta^{N_j} \pi_{-j}^0 \left[\frac{r}{\mu+r} (1-\beta) N_j + \frac{\mu}{\mu+r} (1-\beta^{N_j}) + 1 \right] \\ &< \pi_j^0 + 2\kappa N_j \beta^{N_j} \pi_{-j}^0 \end{aligned} \tag{17}$$

Using a similar logic as before, the bound for \underline{w}_j follows directly from (4)

$$\underline{w}_j \geq \frac{r}{r+\mu} \pi_j^0 + \frac{\mu}{r+\mu} (1-\beta^{N-j}) \pi_j^0 \geq \left(1 - \frac{\mu}{r+\mu} \beta^{N-j} \right) \pi_j^0. \tag{18}$$

Combining (17) and (18) yields

$$\hat{w}_j - \underline{w}_j \leq \left(2\kappa N_j + \frac{\mu}{r+\mu} \right) \beta^{N-j} \pi_j^0$$

Substituting β^{N-j} with the bound from (14), it follows that

$$\hat{w}_j - \underline{w}_j \leq \left(2\kappa N_j + \frac{\mu}{r+\mu} \right) \frac{\mu}{r\Delta_j(\beta)} (\hat{w}_j - \underline{w}_j) \pi_j^0$$

Which implies the desired bound.

$$N_j \geq \frac{r\Delta_j(\beta)}{2\kappa\mu\pi_j^0} - \frac{\mu}{2\kappa(r+\mu)}$$

■

Proof of Lemma 4.3. We know from Lemma 4.2 that, after any such history, $-j$ will mimic a_{-j}^0 for at least $N(\pi^0, \lambda, \mu, r)$ periods. Moreover, if either $2\kappa > 1$ or $\mu < \underline{\mu}(r) = r2\kappa/(1-2\kappa)$, then

$$N(\pi^0, \lambda, \mu, r) \geq \frac{r\Delta_j(\beta)}{2\kappa\mu\pi_j^0} - 1$$

Using (4), j 's continuation value w_j satisfies

$$w_j \geq \left(1 - \beta^{N(\pi^0; \lambda, \mu, r)}\right) \pi_j^0$$

Hence, the difference between π_j^0 and j 's continuation value after a history in which firm $-j$ observes a_j^0 is bounded by

$$\begin{aligned} \pi_j^0 - w_j &\leq \beta^{N(\pi^0; \lambda, \mu, r)} \pi_j^0 < \delta^{N(\pi^0; \lambda, \mu, r)} \pi_j^0 = \frac{1}{\delta} \exp\left((N(\pi^0, \lambda, \mu, r) + 1) \log \delta\right) \pi_j^0 \\ &= \frac{1}{\delta} \exp\left(\log \delta \frac{r \Delta_j(\beta)}{2\kappa\mu\pi_j^0}\right) \pi_j^0 < \frac{1}{\delta} \exp\left(\log \delta \frac{r \Delta_j(\beta)}{2\kappa\mu\hat{\pi}}\right) \hat{\pi} \\ &= \frac{\hat{\pi}}{\delta} \exp\left(-C_1(\lambda, r) \frac{\Delta_j(\beta)}{\mu}\right), \end{aligned}$$

where $C_1(\lambda, r) = r \log \delta / 2\kappa\hat{\pi}$. ■

Proof of Lemma 4.4. To simplify the exposition, I will use the notation $\epsilon_\mu = \epsilon(\lambda, \mu, r)$ in this proof. Suppose j observes a_{-j}^0 at time $t = 0$. For now, condition everything on j never having another revision, a similar argument applies to the case where j gets new revision opportunities. Let \tilde{n} be the number of consumers before $-j$'s next revision, distributed as

$$\Pr(\tilde{n} = n) = \left(\frac{\lambda}{\mu + \lambda}\right)^n \left(\frac{\mu}{\mu + \lambda}\right)$$

Also, let w_n be continuation profits if $-j$ gets a revision after n consumers, and let $\mathcal{N}(\epsilon)$ be the values of n for which the joint continuation profits are ϵ -close to the joint monopolistic profits, formally,

$$\mathcal{N}(\epsilon) = \{n \mid \bar{w}_n \geq (1 - \epsilon)\bar{\pi}^*\}$$

Note that the we can write

$$\begin{aligned} \Pr\left(\bar{w}_{\tilde{n}} \geq (1 - \epsilon)\bar{\pi}^*\right) &= \Pr(\tilde{n} \in \mathcal{N}(\epsilon)) \\ &= \frac{\mu}{\mu + \lambda} \sum_{n \in \mathcal{N}(\epsilon)} \left(\frac{\lambda}{\mu + \lambda}\right)^n = (1 - \gamma) \sum_{n \in \mathcal{N}(\epsilon)} \gamma^n. \end{aligned} \quad (19)$$

where $\gamma = \lambda/(\lambda + \mu)$. I will show that, in any equilibrium and for any $m \in \mathbb{N}$ we have that $\Pr\left(\bar{w}_{\tilde{n}} \geq (1 - m\epsilon_\mu)\bar{\pi}^*\right) \geq m/(m + 1)$.

To see this consider any equilibrium, and let a_j^0 be an algorithm that firm j chooses with positive probability at the current revision. Instead, firm j could choose the alternative algorithm a'_j described as follows.

- First, let (p_n^0) be the sequence of prices induced by a^0 . The algorithm a'_j mimics a_j^0 along this sequence.
- For every n let a_j^n be the continuation algorithm for a'_j after n consumers have arrived, Best responses are unique except in uninteresting cases so let π'_n be the

profit obtained from \hat{a}_{jn}^0 and $-j$'s best response.

- Let $\mathcal{M}(\epsilon) = \{n | \bar{\pi}'_n \in \delta\Pi\}$.
- For $n \notin \mathcal{M}(\epsilon)$, there exists p'_n such that $\pi(p'_n) \gg \bar{\pi}'_n + u\epsilon$
- Firm j appends a'_j offering a grim trigger for p'_n after n for $n \notin \mathcal{M}(\epsilon)$

Next step is to show that best responses to a'_j are weakly better than best responses to a_j^0 and are strictly better for $n \notin \mathcal{M}(\epsilon_\mu)$.

From Lemma (4.3), for $n \in \mathcal{N}(m\epsilon_\mu)$ we have a loss of at most ϵ_μ , and for $n \notin \mathcal{N}(m\epsilon_\mu)$ there is a gain of at least $m\epsilon_\mu$. Hence, using equation (4), we have that

$$\begin{aligned} v'_j - v_j^0 &> \sum_{\tilde{\mathcal{N}}(m\epsilon_\mu)} \beta^m m\epsilon_\mu - \sum_{\mathcal{N}(\epsilon_\mu)} \beta^m \epsilon_\mu \\ &= \left[\frac{1}{1-\beta} - \sum_{\mathcal{N}(m\epsilon_\mu)} \beta^m \right] m\epsilon_\mu - \sum_{\mathcal{N}(m\epsilon'(\mu))} \beta^m \epsilon'(\mu) \\ &= \frac{m\epsilon_\mu}{1-\beta} - \epsilon_\mu(m+1) \sum_{\mathcal{N}(m\epsilon_\mu)} \beta^m \end{aligned}$$

The deviation would be whenever $v'_j - v_j^0 >$, hence, in equilibrium we must have

$$(1-\beta) \sum_{\mathcal{N}(m\epsilon_\mu)} \beta^m \geq \frac{m\epsilon_\mu}{m\epsilon_\mu + \epsilon_\mu} = \frac{m}{1+m} \quad (20)$$

Since $0 < \gamma < \beta < 1$, combining (19) and (20) yields the desired result. \blacksquare

Proof of Theorem 4.1. The proof requires some preliminaries. First, I will use the notation $\nu = \eta\bar{\pi}^M$ and $\gamma := \lambda/(\lambda + 2\mu + r)$. Fix any positive number $\vartheta > 0$ and any $m > 0$, and let

$$\Delta := \frac{1}{C_1(r)} \left(\frac{\log \gamma \log \vartheta}{\log(\lambda + 2\mu) - \log \lambda} - \log(\delta\eta/m\hat{\pi}) \right) \mu,$$

and

$$\epsilon := \epsilon(\Delta; \lambda, \mu, r) = \frac{1}{\delta} \hat{\pi} \exp\left(-C_1(r) \frac{\Delta}{\mu}\right).$$

After some algebra, we have that

$$\epsilon = \frac{\eta}{m} \exp\left(-\frac{\log \gamma \log \eta}{\log(\lambda + 2\mu) - \log \lambda}\right).$$

Now, suppose that we start from a history h in which $\bar{v}_0 \geq (1 - m\epsilon)\bar{\pi}^M$, and let \bar{w}' denote the joint continuation profits after n_0 new costumers have arrived for some arbitrary $n_0 \in \mathbb{N}$. I want to show that

$$\lim_{\mu \rightarrow 0} \inf_{s^* \in S^*} \Pr_{s^*}(\bar{w}' < (1 - \eta)\bar{\pi}^M) = 0$$

From Lemma 4.4 and the missing uniformity step, we know that the joint continuation values after each subsequent revision being greater than $m\epsilon$ is weakly greater than $m/(1+m)$. Moreover, if μ is large enough, $m\epsilon < \eta\bar{\pi}^M$.

Suppose that the joint continuation values since the last revision at n_0 are greater than $(1-m\epsilon)\bar{\pi}^M$. Let \bar{v}_n be the continuation value after n consumers have arrived since the last revision. From (2), we can write

$$\bar{v}_n = \frac{r}{\lambda + 2\mu + r} \bar{\pi}_1 + \frac{\lambda}{\lambda + 2\mu + r} \bar{v}_1 + \frac{2\mu}{\lambda + 2\mu + r} \bar{w}.$$

and thus

$$\begin{aligned} \bar{v}_1 &= \frac{1}{\lambda} \left[(\lambda + 2\mu + r) \bar{v}_n - 2\mu \bar{w} - r \bar{\pi}_1 \right] \\ &\geq \frac{1}{\lambda} \left[(\lambda + 2\mu + r) (1 - m\epsilon(\mu)) \bar{\pi}^M - (2\mu + r) \bar{\pi}^M \right] = \left(1 - \frac{m\epsilon(\mu)}{\gamma} \right) \bar{\pi}^M. \end{aligned}$$

In this case, a sufficient condition $\bar{w}' \geq (1-\eta)\bar{\pi}^M$ is that

$$\frac{m\epsilon\mu}{\gamma^n} \leq \eta \quad \Leftrightarrow \quad n \leq -\frac{\log \eta - \log(m\epsilon)}{\log \gamma}$$

Let \tilde{n} be the number of consumers between revisions. The probability that $\tilde{n} < n$, is one minus the probability that the first n events are consumer arrivals, hence

$$\Pr \left(\tilde{n} < \frac{\log \eta - \log(m\epsilon\mu)}{-\log \gamma} \right) = 1 - \left(\frac{\lambda}{\lambda + 2\mu} \right)^{-(\log \eta - \log(m\epsilon\mu))/\log \gamma}$$

Suppose we want this probability to be less than ϑ , this is equivalent to

$$\epsilon_\mu \leq \frac{\eta}{m} \exp \left(-\frac{\log \gamma \log \vartheta}{\log(\lambda + 2\mu) - \log \lambda} \right),$$

which is satisfied by our previous condition. Hence, I have established that

$$\lim_{\mu \rightarrow 0} \inf_{s^* \in S^*} \Pr_{s^*} (\bar{w}' < (1-\eta)\bar{\pi}^M) \leq \vartheta + \frac{1}{m+1}.$$

Since ϑ and m where arbitrary, we can conclude that

$$\lim_{\mu \rightarrow 0} \inf_{s^* \in S^*} \Pr_{s^*} (\bar{w}' < (1-\eta)\bar{\pi}^M) = 0.$$

■

C. Proof of additional results

C.1. Asymmetry and Leadership

In order to prove Proposition 5.2, it is convenient to introduce some notation. First, for a sequence of prices (p_n) and number x define $\Delta_2((p_n), x)$ as

$$\Delta_2(x) := \inf_{n \geq 0} \left\{ (1 - \delta) \sum_{k=0}^{\infty} \delta^k \pi_2(p_k) - \left[(1 - \delta) \hat{\pi}_2(p_n) + \delta \underline{\pi} \right] \right\}.$$

Also, given a profile of algorithms a inducing a sequence of prices (p_n) , let $\pi(a, n)$ denote the expected discounted profit generated by the first n consumers to arrive, that is,

$$\pi(a, n) = (1 - \delta) \sum_{k=0}^n \delta^k \pi(p_k).$$

Clearly, $\lim_{n \rightarrow \infty} \|\pi(a) - \pi(a, n)\| = 0$. Moreover, the convergence is uniform over a since

$$\pi_j(a) - \pi_j(a, n) = \sum_{k=n+1}^{\infty} \pi_j(p_k) \leq \delta^{n+1} \hat{\pi}. \quad (21)$$

Proof of Proposition 5.2. Let a^S be an algorithm profile that achieves the Stackelberg outcome. Without loss of generality, assume that a^S induces a sequence (p_n^S) along the equilibrium path and deviates after any observable deviation to “always zero”. Since a^S implements the Stackelberg outcome as an equilibrium, it must be the case that $\Delta_2((p_n^S), \delta) \geq 0$.

Now, for some $n_0 \in \mathbb{N}$ and let a' mimic a^S for the first n periods, and then chooses the price profile that yields $(0, \hat{\pi})$. Clearly, regardless of the value of n_0 we have $\Delta_2(a', \delta) > 0$. Moreover, since a' mimics a^S for the first n_0 consumers to arrive, we know from (21) that

$$\pi_1(a') \geq \pi_1(a^S) - \delta^{n_0} \hat{\pi}.$$

Repeating the steps from Lemma 4.2, we know that, after observing a'_1 , form 2 will mimic a'_{-j} for at least n periods, where n is the largest integer satisfying

$$n < N_2(\mu_1) := \frac{1}{\log(\beta_1)} \log \left(\frac{\hat{\pi} \mu_1}{r \Delta_2(a', \beta_1)} \right).$$

It is easy to see that, holding other parameters constant, $\lim_{\mu_1 \rightarrow 0} N_2(\mu_1) = \infty$. Hence, as for μ_1 small enough so that $N_2(\mu_1)$ we have that, by choosing a'_j at every revision opportunity, it guarantees a the followign loer bound for his equilibrium profits

$$\begin{aligned} v_1(\lambda, r, \mu_1, \mu_2) &\geq \frac{\mu_2}{\mu_2 + r} \left(1 - \delta^{N_2(\mu_1)} \right) \pi_1(a') \\ &\geq \frac{\mu_2}{\mu_2 + r} \left(1 - \left(\frac{\mu_1 \hat{\pi}}{r \Delta_2(\beta_1)} \right)^{\log \delta / \log \beta_1} \right) (\pi_1^S - \delta^{n_0} \hat{\pi}). \end{aligned}$$

The result obtains taking the limits as n_0 goes to ∞ . ■

C.2. Collision between impatient firms

Proof of Proposition 5.3. Fix r and λ and any monotone sequence (μ_n) with $\lim \mu_n = \infty$. Let $\underline{v}_n = \sup\{v_1 \mid v_1 = v_2 \wedge v \in V(r, \lambda\mu_n)\}$ be the worst symmetric equilibrium profit given μ_n . We want to show that $\underline{v} := \lim \underline{v}_n = \pi^M$. Feasibility implies that $\underline{v} \leq \pi^M$. Hence, we only have to show $\underline{v} \geq \pi^M$.

Suppose towards a contradiction that $\underline{v} < \pi^M$. By construction, there exists a sequence of strategy profiles $s_n^* \in S^*(r, \lambda\mu_n)$ such that $v(s_n^*) = \underline{v} + \epsilon_n$ for some sequence (ϵ_n) such that $\lim \epsilon_n = 0$. Let w_n^* be the corresponding equilibrium continuation values after any history other than the initial history. From (4), for any algorithm $a_j \in A$ we have that the profit for j from choosing a_j at period 0 if all other choices are made according to s^* can be written as

$$v_{jn}(a_j, s_n^*) = \frac{r}{2\mu_n + r} \tilde{\pi}(a_j, s_{-jn}^*) + \frac{\mu_n}{2\mu_n + r} [\tilde{w}_{n1j}^*(a_j) + \tilde{w}_{n2j}^*(a_j)] \leq \underline{v} + \epsilon_n,$$

where the inequality obtains because s_n^* is an equilibrium.

Now let s'_n be the sequence of strategies described as follows. At date $t = 0$ and along the equilibrium path, both firms always choose the “*always monopolistic*” algorithm denoted by a_j^M (which has a single state ad is thus always feasible). After any history in which a firm has a revision and firm j is revealed to be using an algorithm $a_j \neq a_j^M$, firms use the continuation strategies from s_j^* that result in the continuation values $\tilde{w}_{n1j}^*(a_j)$ and $\tilde{w}_{n2j}^*(a_j)$. Note that s' will be an equilibrium as long as

$$\pi^M \geq \frac{r}{2\mu_n + r} \pi_j(a_j, a_{-j}^M) + \frac{\mu_n}{2\mu_n + r} [\tilde{w}_{n1j}^*(a_j) + \tilde{w}_{n2j}^*(a_j)].$$

Note however that the right hand side of the inequality satisfies

$$\begin{aligned} \text{RHS} &= \frac{r}{2\mu_n + r} \pi_j(a_j, a_{-j}^M) + \left[v_{jn}(a_j, s_n^*) - \frac{r}{2\mu_n + r} \tilde{\pi}(a_j, s_{-jn}^*) \right] \\ &\geq \underline{v} + \epsilon_n + \frac{r}{2\mu_n + r} [\pi_j(a_j, a_{-j}^M) - \tilde{\pi}(a_j, s_{-jn}^*)] \\ &\geq \underline{v} + \epsilon_n + \frac{r}{2\mu_n + r} 2\hat{\pi} \xrightarrow{n \rightarrow \infty} \underline{v} \end{aligned}$$

The assumption $\underline{v} < \pi^M$ thus implies that there exists some N such that s'_n is an equilibrium for $n \geq N$. However, this means that $\underline{v}_n \geq \pi^M$ for $n \geq N$, and, consequently, $\pi^M \geq \pi^M$. \blacksquare